

Are There Decisions Computers Should Never Make?

James H. Moor

The possibility may seem exhilarating or it may seem repugnant, but the possibility should be carefully considered. The possibility is that computers may someday (and perhaps to a limited extent already do) serve not merely as tools for calculation or consultation but as full-fledged decision makers on important matters involving human welfare. In examining this possibility I hope to avoid computerphilia and computerphobia and argue for an empirical approach as a significant component in our assessment of computer activity and its effects. I wish to focus on the issue of decision making because it is in this area that computers have the greatest potential for influencing and controlling our lives. In determining what limits, if any, we should place on the use of computers, we must consider whether there are decisions computers should never make.

DO COMPUTERS MAKE DECISIONS?

It can be objected that asking whether there are decisions computers should never make begs an important question, i.e., whether computers are the sort of thing which can make decisions at all. Before considering this objection, it is useful to distinguish between two senses of making a decision. In the *narrow* sense 'making a decision' refers to the arrival at a decision, i.e., to the selection of a course of action. Processes leading up to the decision are ignored. For example, if one is asked to pick any card during a card trick, then simply selecting a card constitutes making a decision. In the *broad* sense 'making a decision' refers not only to the decision but to processes leading up to the decision as well. Thus, in the broad sense making a decision may involve investigating possible courses of action, evaluating alternative strategies and selecting a course of action based on this investigation and evaluation. For example, in playing checkers one makes a decision by considering various possible moves, weighing the advantages and disadvantages of each and finally selecting a move based on this analysis.

Now, the objection above can be put more precisely. Computers might make decisions (or at least be used to make decisions) in the narrow sense of the term, but computers are not the sort of thing which can make decisions

[217]

in the broad sense. In other words, computers might make decisions in the sense that one can flip a coin to make decisions, but computers are not the sort of thing which can investigate and evaluate alternative strategies in order to select a course of action.

I believe this objection is mistaken. Perhaps its initial plausibility stems from understanding a computer simply as a calculator of arithmetic operations. However, computer activity can be understood in many other ways.¹ One of the most common ways of understanding computer activity is in terms of the execution of an ordinary computer program. In describing this activity, programmers often use decision making language. For instance, a programmer might say that at a certain point in the execution of a program the computer decides whether an inputted string of characters matches another string of characters. Of course, such uses of 'decides' and its cognates might be discounted as nothing more than technical jargon. But there are other situations in which computer activity can be understood as a complex analysis of information resulting in the selection of a course of action. In such cases, decision making language often has a very natural application.

As an example, consider A. L. Samuel's now classic program for playing checkers.² The computer using Samuel's program not only plays checkers, but improves its game with experience. The computer understood as a checker player is naturally described as a decision maker. When its turn comes, the computer must decide what move to make. Moreover, if the computer is to play checkers well, it must base its decisions upon sophisticated decision making processes. As Samuel points out, "There exists no known algorithm which will guarantee a win or a draw in checkers, and the complete explorations of every possible path through a checker game would involve perhaps 10^{40} choices of moves which, at 3 choices per millimicrosecond, would still take 10^{21} centuries to consider."³ The computer, using Samuel's program, makes its decision not unlike human players in that it looks ahead a few moves and evaluates the possible resulting board positions, but it differs from the human decision maker in the manner in which it evaluates board positions. The computer evaluates a board position in terms of a polynomial each term of which represents a parameter of the game, i.e., some configuration of pieces and squares. By playing lots of games, sometimes with itself as an opponent, the computer learns which parameters are important. By altering the weights of the parameters and trying different parameters from a stockpile of them, the computer's evaluation mechanism becomes better and better. Although the computer played poor checkers initially, through competition the computer's abilities improved to the point that it beat a human checker playing champion.⁴

The fact that the computer uses a polynomial to determine its selection of moves does not show that the computer is not a decision maker, for a

human player could make his decisions in the same manner though certainly not as quickly. Indeed, a computer playing checkers is a very clear illustration of a computer making a decision in the broad sense. The computer must analyze the situation, discover what courses of action are available, evaluate the options, and select a course of action based on its information. This is a paradigm of decision making.

One might attempt to buttress the original objection that computers cannot make decisions by assuming that decision making must be done consciously. Certainly we are conscious of much of our decision making, but it is important to realize that we are not conscious of much of it as well. For example, each of us often decides what food to eat or which clothes to wear without being even slightly aware of why a particular decision is made. Much money is invested in marketing research to discover those factors, those "hidden persuaders" as Vance Packard once called them, which can affect our decision making without our being aware of them. Sometimes we can make decisions without even being aware that decisions are being made. For instance, unless we happen to reflect on the situation later, we can make many complex driving decisions in heavy traffic, perhaps while thinking about something else, without being conscious of our own decision making (in the narrow or the broad sense). Since our consciousness of our own decision making can vary from being much aware to completely unaware, consciousness of decision making should not be regarded as an essential feature of decision making.

Finally, it might be argued that it is not computers which make decisions but rather humans who *use* computers to make decisions. But, this point confuses the *power* to make decisions with the *ability* to make decisions. The power to make decisions involves being in the appropriate situation and having the authority to make decisions. For instance, at any time, only one person has the power to make United States presidential decisions although many people may have the ability to make such decisions. The source of this power comes from an election by the people under the Constitution. The fact that we use the president to make decisions is compatible with the president being a decision maker. Similarly, we can delegate decision making power to computers, and the fact that we use computers in this way is compatible with computers being decision makers.

I believe it is important to understand computer activity in some contexts as decision making not only because it is so, but because to see it otherwise tends to minimize our appreciation for the potential impact of computers on our society. To delegate decision making power is to delegate control. Ultimately, the issue is what aspects of our lives, if any, computers should control.

HOW COMPETENT CAN COMPUTER DECISION MAKING BE?

If one grants that at least in principle computers are able to make decisions, it remains a question what kinds of decisions computers can make competently. Since computers are not limited to making random, fixed, or arbitrary decisions, as the checker playing computer illustrates, it may seem that there are no limits to computer decision making. But, the results of logic clearly indicate some limitations. If one accepts Church's thesis that algorithmic computability of a function is equivalent to Turing machine computability of it, then limits of Turing machines are limits of computers. Specifically, the results of the halting problem show that there are decisions even universal Turing machines cannot make effectively, viz., there is no universal Turing machine which can decide for every Turing machine whether or not it will halt. The trouble with this type of limitation is that it seems to apply to humans as well as to computers. Moreover, if one were to seriously set out to decide whether or not sufficiently complex Turing machines would halt, computers, though not infallible, would likely be better at the job than humans.

Therefore, the issue is not whether there are some limitations to computer decision making but how well computer decision making compares with human decision making. In order to make the matter most interesting, I will limit the class of computers to those sorts of electronic and mechanical devices which are ordinarily considered to be computers, i.e., for the purposes of this paper, I wish to specifically rule out considering human beings as computers and considering computers as persons.⁵ Are there, then, decisions which (nonperson) computers could never make as well as humans?

I believe the simple, honest answer is that nobody really knows whether computers can possibly match or exceed human ability at decision making. I wish to advocate an empiricist's position on the question of computer decision making and on the question of computer intellectual abilities in general. My claim is:

- (1) It is essentially an empirical matter what a computer's level of ability is for a given intellectual activity.
- (2) It is possible to gather evidence to determine a computer's level of ability for a given intellectual activity.
- (3) For most kinds of intellectual activities it is still unknown whether or not computers will one day match or exceed human levels of ability.

As a corollary of my general empiricist's position, I want to maintain that for most kinds of decision making, it is still a very open empirical question whether computers will ever have levels of ability which match or exceed human levels. I regard my view as nothing more than common sense, but common sense seems to be somewhat uncommon on this matter. For

instance, some would challenge my view on the grounds that it is not an empirical matter at all. With regard to decision making, I have already responded to this kind of objection. Others who grant that there is an empirical component involved often suggest that the matter is already settled. For instance, in 1958 Herbert Simon and Allen Newell, prominent artificial intelligence researchers, asserted that "there are now in the world machines that think, that learn and that create. Moreover, their ability to do things is going to increase rapidly until—in the visible future—the range of problems they can handle will be coextensive with the range to which the human mind has been applied."⁶ Artificial intelligence workers have clearly demonstrated that computers can possess certain kinds of intellectual abilities. To an amazing extent, today's computers can solve problems, recognize patterns, play games, prove theorems, and use natural language.⁷ Nevertheless, it is just a brute fact that computers do not now possess anywhere near the general intelligence of an average human being, and there is no strong evidence that in the visible future the range of problems they will be able to handle will be coextensive with the range to which the human mind has been applied. The enthusiasm of artificial intelligence researchers for their work is commendable, but at this time the results of their labors do not establish that computers will one day match or exceed human levels of ability for most kinds of intellectual activities.

Some of the critics of artificial intelligence research would also disagree with my view. Hubert Dreyfus concludes his analysis of such research by stating: "Thus, insofar as the question whether artificial intelligence is possible is an empirical question, the answer seems to be that further significant progress in Cognitive Simulation or in Artificial Intelligence is extremely unlikely."⁸ Dreyfus appeals both to the fact that work in these areas sometimes fails and to a phenomenological analysis which he takes to show that there are nonprogrammable human capacities involved in all forms of intelligent behavior.⁹ The argument about failures in early endeavors is not very persuasive since it can be launched against any science in its early stages. The more interesting argument is his phenomenological appraisal which emphasizes, I think correctly, that when we engage in perceptual and intellectual activities, we usually have a global recognition of the situation and can pull out the essential features even in ambiguous contexts. For instance, if I utter the sentence 'Christopher Columbus had global recognition which no computer has ever had' we can all immediately grasp several meanings of the sentence and know which are related to this discussion and which are puns; and yet, we are not aware of any extensive analysis leading up to this understanding. Computers clearly lack such facility with language. Although today's computers can handle some perceptual and linguistic ambiguities, on the whole computers are very much inferior to people on such matters. Computers are not good punsters. Nevertheless, these phenomenological and factual points are not

adequate to establish Dreyfus' conclusion that there are nonprogrammable capacities involved in all forms of intelligent behavior. What appears unlawful, even capricious, at one level may be perfectly lawful at another. It remains a possibility that activities of which we are not aware, but which underlie intelligent behavior, can be expressed in terms of computable functions. If this is the case, then computers might one day carry them out.

A task which may seem unprogrammable is the selection of appropriate hypotheses in science. But for certain families of molecules a computer using the program DENDRAL is an expert in identifying the molecular structure which best explains data produced by mass spectrometers. The computer is an expert in doing this even when compared with the best human performance.¹⁰ On the other hand, for most of chemistry the computer's performance in selecting appropriate hypotheses is at the novice level or worse. The point is that there is enough evidence from artificial intelligence research to be suspicious of dogmatic claims that computers will never be able to accomplish certain feats of intelligence; yet, today there is not nearly enough evidence to support the conclusion that computers will someday match or exceed human intellectual ability in general or human decision making ability in particular.

HOW CAN COMPUTER DECISION MAKING COMPETENCE BE JUDGED?

Empirical investigation will allow us to refine our judgments about the nature of computer abilities, but what sort of evidence counts? Since competence is the ability to perform at a given level of accomplishment, obviously the computer's performance will be one important source of evidence in evaluating competence. With regard to decision making, two features of the performance are relevant: (i) the decision making record and (ii) the justifications offered for the decisions. These two types of evidence will carry various weights depending upon the kind of decision making and the circumstances in which it is made.

Specifically, I wish to distinguish two extremes of decision making: decision making under clear standards and decision making under fuzzy standards. Many cases of decision making lie between these two extremes. In making decisions under clear standards every decision (or series of decisions) can be clearly classified as either correct or incorrect. For example, deciding which horse to bet on to win a race is a decision made under clear standards. The horse will either win or not win. After a number of such decisions, preferably made in a variety of situations, the decision maker will have established a clear record of correct decisions vs. incorrect decisions. In the case of decision making under clear standards the justifications offered for decisions are usually not very important in evaluating competence. If the bettor routinely picks winning horses, it hardly matters if he is unable to produce justifications for his decisions. If he cannot pick winning horses, justifications for decisions are small consolation. However,

if the decision making record is not available or for some reason is not trusted, then certainly the justifications given for the decisions can be important in evaluating decision making under clear standards.

In decision making under fuzzy standards at least some of the possible decisions (or series of decisions) will be difficult to classify as correct or incorrect. For many people deciding which career to pursue is an example of decision making under fuzzy standards. The fuzzier the standards are the more difficult it is to establish a clear record of correct decisions vs. incorrect decisions, and the more important it is to provide some justifications for the decisions. For instance, it may be impossible to determine how many 'correct' and 'incorrect' grades a professor gives, but one can evaluate his decision making by checking his justifications for assigning individual grades.

The checker playing computer is making decisions under clear standards in that a series of decisions about moves either leads to a win or it does not. It is impressive that the computer has beaten a checker playing champion; but in order to really establish its competence, the computer would have to establish a substantial record of play against a variety of opponents.

Some computers can make decisions under fuzzy standards and can offer justifications for their decisions. For instance, consider MYCIN, an interactive program that "uses the clinical decision criteria of experts to advise physicians who request advice regarding the selection of appropriate antimicrobial therapy for hospital patients with bacterial infection."¹¹ Advice of this kind has practical importance because often drug therapy has to be recommended before a positive identification of the bacteria can be made and because not every physician is a specialist in the subject. The computer asks the physician for information about the situation including the results of laboratory tests. The computer will give its conclusions about the identities of the organisms; and upon asking a few more questions about the patient's allergies, renal and hepatic status, the site(s) of the infection, etc., the computer will formulate and recommend therapy. MYCIN contains a set of over 200 production rules each of which state a set of preconditions and a conclusion or action to be taken if the preconditions occur. The production rules not only give the computer a basis for decisions, but allows it to offer justifications for its decisions. The computer can explain either why certain information is important in terms of its goals or how it arrives at its conclusions. Thus, a physician has a check on the computer's competence without knowing the computer's overall decision making record. The decision making record cannot be completely clear since even experts disagree about what is a 'correct' or on 'incorrect' decision in some cases, e.g., in recommending certain therapies. Nevertheless, in a preliminary evaluation, MYCIN's therapy recommendations were acceptable to the experts in 75% of the cases.¹² The workers on the MYCIN project hope to increase the computer's competence in this type of decision making

and to extend the production rule methodology to other areas.¹³ But whatever abilities computers may eventually acquire using the production rule approach, the MYCIN program illustrates that a computer's performance can be such that one can evaluate the computer's competence even in decision making under fuzzy standards.

Another kind of evidence about a computer's competence in decision making results from an analysis of the internal operation of the computer, perhaps in terms of a computer program. This kind of evidence is not essential, at least not in principle; for if the performance is good enough, it will provide a sufficient basis for a justified inductive inference about the computer's competence.¹⁴ We often infer that other humans are competent decision makers without having any information about their internal operation except indirectly through performance. Nevertheless, as a practical matter it certainly can be very useful to have such information, e.g., in those situations in which a justification for a decision is an important piece of evidence but which the computer can not provide as part of its performance. Obviously, in the development stage one must pay close attention to the program for even the most novice programmer knows that a well-thought-out program may not result in the performance expected and a performance which is good in general may be the result of a program with hidden "bugs."

Usually it is not very helpful in assessing a computer's competence at decision making to simply ask whether the computer makes its decisions on the same basis (in the sense of internal operation) that humans do. The answer can almost always be "yes" or "no" depending upon how the activity is described. Yes, the computer checker player is like a human checker player in that it looks ahead a few moves. Or no, the computer checker player is not like a human checker player in that it uses a polynomial to evaluate board positions.¹⁵ What is crucial is that the basis be capable of reliably generating a reasonable level of performance (including producing justifications when relevant). Indeed, if a computer does exceed human competence in decision making, it is very likely the basis for its decisions will be different from the human basis.

ARE THERE DECISIONS COMPUTERS SHOULD NEVER MAKE?

The empirical position I am advocating undercuts a lot of argumentation about which decisions computers should and should not make. Joseph Weizenbaum states:

What could be more obvious than the fact that, whatever intelligence a computer can muster, however it may be acquired, it must always and necessarily be absolutely alien to any and all authentic human concerns. The very asking of the question, "What does a judge (or a psychiatrist) know that we cannot tell a computer?" is a monstrous obscenity. That it has to be put into print at all, even for the purpose of exposing its morbidity, is a sign of the madness of our times.

Computers can make judicial decisions, computers can make psychiatric

judgments. They can flip coins in much more sophisticated ways than can the most patient human being. The point is that they ought not be given such tasks. They may even be able to arrive at "correct" decisions in some cases—but always and necessarily on bases no human being should be willing to accept.¹⁶

Weizenbaum claims that computers are outsiders to human affairs just as humans are sometimes outsiders to other human cultures. Outsiders will have bases for decisions which "must be inappropriate to the context in which the decision is to be made."¹⁷ But this argument confuses lack of information with lack of competence. There may be good reasons not to grant outsiders the power to make some decisions, but there is no reason in principle why an *informed* outsider cannot be a competent decision maker. A physician who is an outsider may be more competent to make medical decisions than anybody in a primitive tribe. A computer which never has a bacterial infection may be very competent in making decisions about them.

Weizenbaum does not make it clear whether by "bases" he means the internal operation of the computer or the sorts of justifications it could give for making its decisions. But for neither case has he demonstrated that they must always and necessarily be such that human beings should not accept them. Weizenbaum's examples—judicial decisions and psychiatric judgments—are cases of decision making under fuzzy standards. It is possible to evaluate a computer's competence in these areas by paying close attention to the sorts of justifications the computer gives for its decisions. It is at least conceivable that the computer might give outstanding justifications for its decisions ranging from detailed legal precedents to a superb philosophical theory of justice or from instructive clinical observations to an improved theory of mental illness so that the competence of the computer in such decision making was considered to be as good or better than the competence of human experts. Empirically this may never happen but it is not a necessary truth that it will not.

Perhaps more importantly, an empirical attitude challenges an uncritical acceptance of computer competence. It is far too easy to try to justify a decision by simply saying "the computer says so." Such a reply should carry no weight unless the computer's competence has been rigorously tested. It is always relevant to raise two competency questions when computer decision making occurs—"What is the nature of the computer's (alleged) competence?" and "How has the competence been demonstrated?" A company spokesman might announce that a computer has decided there should be a 20% layoff when in fact the computer has done nothing more than determine which 20% of the firm's employees has least seniority. The problem is not just that one group might deceive another about the computer's competency but that even immediate users of the computer may take the computer's word too uncritically. In a nuclear age in which some of the decision making about whether to launch missiles is in part

made by computers, the possibility of deception about computer competency is a matter of great importance.

Thus, the first step in determining what kinds of decisions computers should and should not make at particular times is to determine what kinds of decisions computers can and cannot make competently at those times. But, there remains a question of values. Even if someday computers are competent to make a wide range of important decisions, should certain kinds of decision making be forbidden to computers? I believe that the proper answer suggests itself when considering why the following three maxims, though initially plausible, are really unsatisfactory.

DUBIOUS MAXIM #1 *Computers should never make any decisions which humans want to make.* This is a somewhat plausible maxim since we obviously enjoy the pleasure and freedom involved in making many of the decisions which affect our lives. A computer could competently decide which shoe a person should put on first in the morning, but clearly such a meaningless intrusion into a person's affairs would greatly reduce the quality of his life. However, this maxim is unsatisfactory in general because there can be other factors which outweigh the benefits of the freedom and pleasure humans derive from doing the decision making. For example, even if humans would like to make certain medical decisions, it might be the case that a computer existed which could make them far better. If the computer's diagnosis and suggestions for treatment would result in a significant savings of lives and reduction of suffering compared with human decision making on the subject, then there is a powerful moral argument for letting computers decide.

DUBIOUS MAXIM #2 *Computers should never make any decisions which humans can make more competently.* This also seems like a very reasonable maxim. We do not want the computer to make life-or-death medical decisions if the computer is less competent than human decision makers. But again the maxim is too limited because it neglects other considerations. Some activities, e.g., certain kinds of factory work or prolonged space travel, may be so boring, time-consuming, or dangerous that it would be morally better to use computers, even if this involved sacrificing some competency in decision making, in order to spare humans from enduring such experiences.

DUBIOUS MAXIM #3 *Computers should never make any decisions which humans cannot override.* This maxim seems most reasonable of all especially if it is set against a background of numerous science fiction tales in which computers take control and humans become their slaves. But there could be situations in which it would be morally better to make it impossible, at least practically speaking, for humans to override computer decisions. Suppose that when people drive cars, tens of thousands of people are killed in automobile accidents, hundreds of thousands are injured, and millions of dollars are lost in property damage. But when computers

drive cars, not only are human transportation needs carried out more efficiently but there is a substantial reduction in deaths, injuries, and property damage. Further suppose in those cases in which humans override computer driving decisions, the accident rate soars. Under such circumstances there is a persuasive moral and prudential argument to have computers do the decision making and not to allow humans to override their decisions.

What I am advocating is that we regard computer decision making instrumentally. For particular situations we must determine whether using computer decision makers will better promote our values and accomplish our goals. The maxims above suggest important considerations but are inadequate as general rules because situations may arise in which the consequences are far better if the maxims are violated. This approach is a natural extension of the empiricist's position described earlier. Within the context of our basic goals and values (and the priorities among them) we must empirically determine not only the competence of the computer decision maker but the consequences of computer decision making as well.

This instrumental view of the value of computer decision making leads to the answer to the question what decisions computers should never make. Computers should never decide what our basic goals and values (and priorities among them) should be. These basic goals and values, such as the promotion of human life and happiness, decrease in suffering, search for truth and understanding, etc., provide us with the ultimate norms for directing and judging actions and decision making. By definition there are not further goals and values by which to evaluate these. Since we want computers to work for our ends, we obviously want to prohibit computers from deciding to change these ultimate norms, e.g., promoting computer welfare at the expense of human welfare or taking inconsistency to be the mark of good reasoning.

To prohibit computers from making decisions about basic goals and values (and the priorities among them) is, of course, not to limit computer decision making very much. Our basic goals and values remain fairly constant and humans rarely decide to change them. Thus, there is a wide range of possible decision making which computers one day might justifiably perform. Nevertheless, I believe there is a very legitimate concern that increased computerization of our society will lead to dehumanization of our lives. The proper root of this concern is not that computers are necessarily incompetent or inherently evil. It may be the case that one day computers will make the major decisions about the operations of our society better than humans with the result that the quality of human life is substantially improved. The root of concern about increased computerization should be focused on the issue of responsibility. By assumption, the kind of computers under discussion are not persons; and although they are causally responsible for their decisions, they are not legally or

morally responsible for their decisions. One cannot sue a computer. Therefore, humans have not only an initial responsibility, but a continuing responsibility to raise the competency and value questions whenever computer decision making is at issue. First, what is the nature of the computer's competency and how has it been demonstrated? Secondly, given our basic goals and values why is it better to use a computer decision maker in a particular situation than a human decision maker? The danger is that our responsibility can be easily undermined by strong pressures, e.g., economic incentives, not to investigate and answer these questions. The dehumanization which results can either be in the form of computers making decisions which humans should make or vice versa. Of course, if the delegation of decision making power is carried out responsibly, we may be creating a much more humane society. Some of the most humanistic decisions may well come from decision makers which are not human.

Dartmouth College

NOTES

1. On one level a computer is nothing more than a physical system and can be explained as such. For the computer to perform even simple arithmetic calculations, we must interpret its activities symbolically. Obviously, there is a wide range of possible interpretations. See James H. Moor, "Explaining Computer Behavior," *Philosophical Studies*, 34 (1978): 325-27.
2. A. L. Samuel, "Some Studies In Machine Learning Using the Game of Checkers," *Computers and Thought*, ed. Edward A. Feigenbaum and Julian Feldman (New York: McGraw-Hill, 1963), pp. 71-105.
3. Samuel, p. 72.
4. Samuel, pp. 103-105.
5. I want to separate the question of computers being persons from the main issue because I believe interesting results follow about computer decision making without raising matters of civil rights. The set of computers I am considering will have members which may be very good at particular kinds of decision making but no one member of the set will have sufficient variety of decision making abilities (among other things) to be considered a person.
6. H. A. Simon and A. Newell, "Heuristic Problem Solving: The Next Advance in Operations Research," *Operations Research* 6 (1958): 8.
7. For a nice summary of artificial intelligence work see Patrick Henry Winston, *Artificial Intelligence* (Reading, Mass.: Addison-Wesley, 1977).
8. Hubert Dreyfus, *What Computers Can't Do* (New York: Harper and Row, 1972), p. 197.
9. Dreyfus also has an argument based on the digital/analogue distinction. See James H. Moor, "Three Myths of Computer Science," *The British Journal of Philosophy of Science*, 29 (1978): 213-22.
10. E. A. Feigenbaum, B. G. Buchanan and J. Lederberg, "On Generality and Problem Solving: A Case Study Using the DENDRAL Program," *Machine Intelligence*, 6, ed. B. Meltzer (Edinburgh: Edinburgh University Press, 1971), p. 165.
11. Edward H. Shortliffe, Randall Davis, Stanton G. Axline, Bruce G. Buchanan, C. Cordell Green, and Stanley N. Cohen, "Computer-Based Consultations in Clinical Therapeutics: Explanation and Rule Acquisition Capabilities of the MYCIN System," *Computers and Biomedical Research* 8 (1975): 303.